

Abstract of the Disclosure

A method learns a structure of a video, in an unsupervised setting, to detect events in the video consistent with the structure. Sets of features are selected from the video. Based on the selected features, a hierarchical statistical model is updated, and an information gain of the hierarchical statistical model is evaluated.

Redundant features are then filtered, and the hierarchical statistical model is updated, based on the filtered features. A Bayesian information criteria is applied to each model and feature set pair, which can then be rank ordered according to the criteria to detect the events in the video.